NAS2-11530

Reliability Mechanisms of the FDDI High Bandwidth Token Ring Protocol

Marjory J. Johnson

May, 1985
(Revised July, 1985)

Research Institute for Advanced Computer Science
NASA Ames Research Center

RIACS TR 85.5

Research Institute for Advanced Computer Science

Reliability Mechanisms of the FDDI High Bandwidth Token Ring Protocol

Marjory J. Johnson

May, 1985

Research Institute for Advanced Computer Science
NASA Ames Research Center

The Fiber Distributed Data Interface (FDDI) is an ANSI draft proposed standard for a 100 megabit per second token ring using fiber optics as the transmission medium. The FDDI design represents an extensive effort to incorporate reliability mechanisms as an integral part of the design. These mechanisms provide fault detection and isolation functions, monitoring functions, and configuration functions. The purpose of this paper is to discuss these reliability mechanisms, and to compare and contrast them with reliability mechanisms which have been incorporated in other local area network architectures.

This work was supported by the National Aeronautics and Space Administration under NASA Contract NAS2-11530 to the Universities Space Research Association (USRA).

Reliability Mechanisms of the FDDI High Bandwidth Token Ring Protocol

Marjory J. Johnson
Research Institute for Advanced Computer Science
NASA Ames Research Center
Moffett Field, California 94035

1. Introduction

The study of reliability reported herein was motivated by a study of reliability needs of the local area network (LAN) for the Space Station. Functions supporting its reliability should be automated, for even though the Space Station will be manned, the crew will be busy with assigned tasks and will have little time for routine computer maintenance and repair.

The Fiber Distributed Data Interface (FDDI) is an ANSI draft proposed standard for a 100 megabit per second fiber-optic token ring. Because it is designed specifically for high bandwidth and for fiber optics, it is an attractive candidate for use on the Space Station. The purpose of this paper is to describe and evaluate the reliability mechanisms that have been incorporated in the FDDI protocol.

One aspect of reliability is transmission management functions, such as flow control, congestion control, buffer management, acknowledgement schemes, and error recovery, all of which are included in the Open Systems Interconnection (OSI) network model. Functions of this type are not discussed further in this paper. Another aspect of network reliability is management of the various OSI layers, something called "computer network control" in [18]. Additional network management functions are necessary for LANs, since their broadcast nature means that activities of a single station (i.e., network interface unit) affect the entire network. These functions, which are not part of the basic OSI model, fall into three categories: fault detection and isolation, monitoring, and configuration [18].* Automation of these functions is desirable, so that they can be accomplished in a manner transparent to the user.

For each of the three network management categories identified above, this paper presents a discussion of the general nature of the functions in that category, followed by a discussion of the specific mechanisms provided by FDDI. The final section compares the reliability mechanisms of FDDI with those in other LAN architectures. Related work includes a paper which discusses error handling in the IEEE 802.4 token-passing bus [16] and a paper which discusses the physical configuration and the reliability of the access protocol of an IBM token ring whose design is similar to that of the IEEE 802.5 token ring [1].

2. FDDI Token Ring

In this section we present a brief description of the FDDI media access protocol. FDDI is one of only a very few LAN access protocols designed specifically

---

*The need for these functions is now widely recognized in the literature. Both IEEE and ANSI are addressing network management functions in their development of LAN standards. The IEEE 802 LAN documents include a general document, 802.1, currently under development, which treats network management as an extension of OSI layer management. The ANSC X3T9.5 committee, the ANSI committee which is developing FDDI, is developing a separate Station Management (SMT) document [7]. Similar to the IEEE work, the X3T9.5 committee is developing Station Management as an extension of OSI layer management. By definition, Station Management is an entity within each station which "monitors station activity and exercises overall control of station activity [17]."

for high bandwidth and for fiber optics. (See [2], [15], and [20] for discussions of others.) It represents the only current effort to develop a standard for a 100 megabit per second fiber-optic LAN. Consequently, the FDDI effort is attracting attention throughout the computer and communications industry.

The FDDI access protocol is based on the IEEE 802.5 token ring, modified to support the higher data rates. (As explained in [11], "There is more to high speed protocol design than taking a low speed protocol and changing the numbers in it.") A key difference between the FDDI and 802.5 Media Access Control (MAC) protocols is the method of passing the token after transmission. The FDDI method is for a station to issue the token immediately after it completes transmission of its last frame, whereas the 802.5 method is for a station to wait until the first frame it sent propagates back around to it before issuing a new token. Other modifications are presented in [11]. The designers originally conceived of FDDI either as a network to interconnect supercomputers or to connect high speed computers with peripheral devices in a computer center environment, where high speed data transfer is essential, or as a backbone network to connect smaller LANs. However, other configurations (including a frontend network) are now considered likely.

The FDDI protocol provides two classes of service, synchronous service and asynchronous service. Synchronous service is for applications, such as real-time control, that require access to the channel within a specified time period. Each station is allotted a percentage of the bandwidth for its synchronous transmission needs. Synchronous frames can be transmitted by a station whenever it receives the token. Asynchronous transmission is permitted only if the token is rotating quickly enough, as indicated by an internal station parameter, called

Late_Ct. Thus, synchronous transmission receives highest priority. Multiple levels of priority may be assigned for asynchronous traffic.

Access to the network and scheduling of the network is controlled by timers that interact in a relatively complex manner. Each station has its own set of timers and its own Late_Ct parameter. As part of the ring initialization process, the stations negotiate a value for the Target Token Rotation Time (TTRT), a parameter which specifies the expected token rotation time. Each station requests a value that is fast enough to support its synchronous traffic needs. The shortest requested time is assigned to T_Opr, and then this value is used to load each station's Token Rotation Timer (TRT). A station's Late_Ct is incremented each time its TRT expires, and it is cleared each time the station receives the token. The token is considered to arrive at a station on time if its TRT has not expired since the station last received the token, i.e., if Late_Ct = 0. Otherwise, the token is considered to be late.

The amount of time that a station may hold the token and transmit frames is governed by its TRT, its Token Holding Timer (THT), its synchronous bandwidth allocation, and its parameter Late_Ct as follows:

1. If the token is on time, i.e., if Late_Ct = 0, then the current value of TRT is placed in THT, and TRT is reset to T_Opr. Both asynchronous and synchronous frames may be transmitted.

2. If the token is late, i.e., if Late_Ct ≠ 0, Late_Ct is cleared, but the TRT is not reset. Only synchronous frames may be transmitted in this case.

3. TRT is enabled during transmission of all frames, synchronous and asynchronous. THT is enabled only during transmission of asynchronous frames.

4. The length of time an individual station may transmit synchronous frames may not exceed its synchronous bandwidth allocation. The length of time that a station may transmit asynchronous frames is governed by its THT; no asynchronous frames may be transmitted after its expiration. No frames of either class may be transmitted after expiration of the station's TRT.

For more information about the protocol see [4], [11], [12], [14], and [17].

3. Fault Detection, Isolation, and Correction Functions

Ideally, failure of any device attached to the network should not cause failure of the entire network, only the function or service offered by that individual device. To facilitate ease of maintenance and repair, the network should be self-diagnosing, i.e., the network should continuously monitor itself to detect faults. When a problem is detected, the network should be able to determine the nature of the problem and to isolate the fault to a single component or to a small group of components.

There are two major types of errors in a LAN, protocol-related errors and errors caused by failure of a physical resource. The FDDI protocol features several mechanisms which automatically detect the presence of errors of either type and which automatically initiate ring recovery. If the error is protocol-related, the ring recovery process will correct it; otherwise, the ring automatically physically reconfigures itself so as to bypass the failed resource.

3.1. Physical Reconfiguration

A major concern with a token ring architecture is that a break in the ring, caused either by a failed link or a failed station, will cause failure of the entire network. FDDI provides a solution to this single-point-of-failure problem by providing a redundant ring and optical station bypasses, as follows.

The FDDI draft proposed standard specifies that the network should consist of both a primary ring and a secondary ring, rotating in opposite directions. The secondary ring may be used only as a backup ring, or transmission may be scheduled for both rings under normal operations, thus doubling the effective transmission rate of the ring. There are two types of stations, Class A stations and Class B stations. Class A stations are attached to both rings. Class A stations may be equipped with an optical bypass, so that they may be switched out of the ring in the event of failure. Class B stations have only one physical connection to the network. Consequently, a Class B station must be connected to the network via a Class A connection, called a concentrator. "Traffic into a Class A node is made to go back and forth to the attached Class B nodes before it flows to the next Class A node further downstream [6]." Class B nodes may be electronically bypassed within the concentrator if the station fails or is switched off. In the event of link failure on the primary ring or the secondary ring or both, the network is automatically reconfigured (actually wrapped around) using a combination of the operational links of both rings. Multiple link failures may divide the original ring into two or more isolated smaller rings. Diagrams in [14] and [17] illustrate how the network reconfigures itself after failure of one or more links connecting Class A stations. Station Management (SMT), the control entity within each station, automatically manages the necessary reconfiguration. Failure of a link which attaches a Class B station to a concentrator is, of course, equivalent to failure of the Class B station.

3.2. Frame Stripping

A commonly cited problem with a token ring access protocol is a frame which circulates forever on the ring. If the token is issued only after the frame is stripped from the ring, normal operation of the ring would be disrupted until the problem was corrected. This will not occur with FDDI, since the token is issued by a station immediately after it finishes transmitting. However, it is still undesirable for a frame to continue circulating around the ring. The destination station would have to handle the duplicate frames, and counters in all the stations would be affected.

The FDDI protocol provides for frame removal in an orderly way. The station which transmits a frame is responsible for stripping it from the ring. Normally, when a station is not transmitting, it repeats whatever it receives on the ring. When it recognizes its own address in the source address field of a frame, it absorbs the rest of the frame and pumps idle symbols onto the ring instead. This actually leaves a truncated frame on the ring, since the fields at the beginning of the frame will have already been repeated. All other stations will recognize the frame as a remnant and will disregard it when they encounter the idle symbols before an ending delimiter, and this remnant will be absorbed from the ring when it encounters a transmitting station. Even if the station which transmits a frame fails before the frame circulates back around to it, the frame will eventually be removed from the ring when it encounters a transmitting station; normal operation of the ring is not disrupted.

3.3. FDDI Timers

Each station on the ring needs to be able to determine the status of the network, i.e., whether or not the ring is operational. Sensing transmission on the line isn't sufficient evidence to conclude that the ring is operating properly, because this transmission could represent random noise or a jabbering transmitter. To indicate proper operation of the ring, the transmission must be meaningful and each station must be given an opportunity to transmit within a reasonable amount of time.

Two timers included in the FDDI MAC protocol, the Valid Transmission Timer (TVX) and the Token Rotation Timer (TRT), are used to maintain proper operation of the ring by detecting problems and by forcing initiation of ring recovery when necessary. Each station has its own TVX and TRT. These timers work as follows.

The Token Rotation Timer's fault detection and recovery functions are closely related to its ring-scheduling function. The TRTs in all the stations work together to ensure that the amount of time required for the token to circulate around the ring is bounded above by twice T_Opr. Expiration of a station's TRT when the token isn't late, i.e., when Late_Ct = 0, is a means of hurrying the token along, by preventing the station from transmitting asynchronous frames when it receives the token. Expiration of a station's TRT when the token is already late, i.e., when Late_Ct $\neq$ 0, indicates a serious ring problem, and forces the station to initiate ring recovery. It is proved in [13] that under normal ring operation, the token is guaranteed to return to a station before its TRT expires a second time. That is, expiration of a station's TRT when Late_Ct $\neq$ 0 will occur only if there is a serious problem and ring recovery is indeed necessary.

Thus, the TRT detects and attempts automatic correction of errors which make the token appear to be lost.

The Valid Transmission Timer maintains proper operation of the ring by ensuring that the transmission on the channel is meaningful. This timer is reset each time that a "nonrestricted" token is received or that a valid frame is received. (Tokens for general operation of the ring are "nonrestricted." Stations may hog the channel through the use of a special mechanism called a "restricted" token.) Expiration of a station's TVX forces the station to initiate ring recovery. The default value of the TVX is large enough so that expiration will only occur if the invalid transmission is of a long-term nature. This means that random noise on the line won't initiate ring recovery.

Station Management may request that the station capture the token before a frame is placed in the transmission queue. Upon capture of the token, the station transmits idle symbols as a filler until it has a frame ready for transmission. To prevent expiration of the TVX during this time, Station Management can send void frames. Void frames are interpreted and counted as valid frames, but they are not copied into MAC's receive buffer. Hence, receipt of a void frame will reset a station's TVX, but it will not contribute any congestion in the individual stations.

The default value of the TVX is much shorter than the default value of the TRT. If the token is garbled, then the TRT would eventually expire. However, the TVX will detect the transmission problem much earlier, thus improving the efficiency of the ring. In addition, the TVX will detect a problem that the TRT will not. When stations are finished using a restricted token to hog the ring, the last one is supposed to issue a nonrestricted token. If it erroneously issues a restricted token, expiration of the TVX prevents this token from circulating forever and effectively blocking access to the ring.

3.4. FDDI Ring Recovery Process

Ring recovery is the process of restoring ring operation by negotiating a TTRT value and by issuing a new token. A unique station is selected in a distributed manner to issue this new token.

Ring recovery is initiated if a serious problem with ring operation is detected. For example, ring recovery is initiated if too long a time passes between successive receipts of valid transmissions (either tokens or frames), if the token is lost, or if there is a jabbering transmitter or a deaf receiver. Recovery will be successful if the problem is related to the access protocol; recovery will not be successful if the problem is caused by a physical break in the ring. Specific mechanisms that trigger the recovery process include expiration of various timers, receipt of frames indicating that a station on the ring has detected abnormal operating conditions, or a SMT command to initiate recovery. Ring recovery will also be required if a station's parameter settings are inconsistent with the operational parameters of the network as a whole, e.g., if a station requires a faster token rotation time to support its synchronous needs than is currently being supported by the network.

The Claim Token Frame is the mechanism by which ring recovery is attempted. Stations on the ring are linearly ordered by the requested TTRT, with ties resolved by station address. The station with highest precedence is the station which requires the fastest token rotation time. If two stations request the same TTRT, then the station with the higher address takes precedence. The claim token arbitration procedure, if it is successful, selects the highest precedence station to issue the new token. It works as follows. When ring recovery is initiated at a station, the station's transmitter continuously sends Claim Token Frames. The information field of these frames contains that station's requested value for the Target Token Rotation Time (TTRT). The station continues to send these Claim Token Frames until either the Claim arbitration is resolved or the process is preempted. If a station receives a Claim Token Frame from a lower precedence station, then it either starts or continues to transmit its own Claim Token Frames. If a station receives a Claim Token Frame from a higher precedence station, then it defers and repeats all subsequent Claim Token Frames. Hence, at most one station (necessarily, the highest precedence station on the ring) will succeed in receiving its own Claim Token Frames. If this happens, the arbitration is resolved, and this highest precedence station issues a token to restart ring operation.

If ring recovery is successful, the ring is restored to normal operation within only a few milliseconds. Ring recovery consists of two components, pinpointing the highest precedence station on the ring and setting station parameters when the new token is issued. Since Claim Token Frames are short, requiring .00256 ms transmission time, and since at most two cycles of sending them would be necessary (depending on the location of the highest precedence station relative to the station which initiated the recovery process), very little time would be required by the first component. Two initial token rotations are then required before normal ring operation is resumed, i.e., before asynchronous transmission is allowed. The purpose of the first token rotation is to set station parameters and to align timers. Only synchronous transmission is allowed during the second token rotation. Hence, the second component of ring recovery also requires little time.

Claim arbitration should be successfully resolved if there is no physical break in the ring, unless the arbitration is interrupted prematurely. Events that would terminate the arbitration process are either a MAC_Reset signal (discussed in Section 5.5) or receipt of a Beacon Frame (described below). Hence, if recovery is initiated because of a token-related problem, such as a lost token, then the ring will be reinitialized automatically and in a distributed manner.

3.5. FDDI Beacon Frame

The Beacon Frame is used to indicate a problem of a disruptive nature, requiring physical reconfiguration of the network. An important difference between the Claim Token Frame and the Beacon Frame is that the source of the problem is pinpointed with the Beacon Frame, whereas it is not with the Claim Frame. This is because the Claim Frame process, if it runs to completion, leaves the highest precedence station on the ring in control. In contrast, any station which receives a Beacon Frame defers and repeats the frames it receives. Thus, the station immediately downstream from the problem will be the only station to persist in sending Beacon Frames, and will be designated in the source address of the frames. Thus, the situation that is being "beaconed" is located in the transmitter of the upstream station, the receiver of the station which is sending the Beacon Frames, or the link in between.

Beacon Frames are transmitted if there is a change in ring topology. When a break in the ring occurs, causing the token to appear to be lost, ring recovery will be attempted first. However, recovery will fail, because no station will be able to receive its own Claim Token Frames. If a station's TRT expires while it is sending its own Claim Token Frames, then that station begins to transmit Beacon Frames. If a station receives its own Beacon Frames, then that station sends Claim Token Frames again to attempt ring recovery. If there is indeed a break in the ring, then no station will receive its own Beacon Frames. Station Management will automatically physically reconfigure the ring to bypass the failure.

Another type of topology change occurs when a Class A station is inserted into the ring. The station insertion process has not been completely developed. The description in the current version of the Station Management document states that after either a Class A or a Class B station is inserted into the ring, it will transmit Beacon Frames. In this way it could announce its existence and identity to all other stations on the ring. The ring recovery procedure would follow, to negotiate TTRT, to set T_Opr values, and to determine which station would issue a token to resume operation of the ring.

Beacon Frames are also transmitted upon command from SMT. SMT continuously monitors network performance. If it detects a significant degradation of performance, SMT can issue a Beacon Frame to purge the ring of all traffic. The beaconing station remains in control of the network and can transmit recovery information to other stations on the ring, using an "immediate" service class, a class to be used only for recovery purposes. When this service class is used, it is not necessary to capture the token before transmitting.

4. Monitoring Functions

It is an "unsettling fact that current performance models do not appear to predict the performance of real-world network implementations [19]." This is because of simplifying assumptions which must often be made in order to obtain analytical results. Thus, the only way to determine the true performance of a network is on-line collection and analysis of statistics about network traffic.

Monitoring functions are typically divided into three areas: performance measurement, performance analysis, and artificial traffic generation [18]. Performance measurement and analysis are necessary to determine both performance of the access protocol and performance of the network in general. Artificial traffic generation is necessary so that network performance can be observed under a controlled load. This can be a useful tool in the laboratory before the network becomes operational. After the network is in place, artificial traffic generation can help to detect potential problems before they become serious, and it can help in planning for future growth of the network.

FDDI monitoring functions monitor performance of the access protocol only. Monitoring of general network performance is considered to be outside the scope of the standard, since it is not necessary for interoperability of stations in the network.

4.1. Performance Measurement

Several mechanisms have been provided by the MAC protocol for the gathering of statistics and the communication of status information from MAC to Station Management. As stated above, the objective of FDDI monitoring is fault detection and isolation, not determining general network performance.

4.1.1. FDDI Frame Status Indicators

Frame status indicators (four bits long) are located at the end of every frame. Error Detected (E), Address Recognized (A), and Frame Copied (C) indicators are mandatory; others may be added if desired. Each station has flags which correspond to the mandatory indicators. These three indicators are initially reset. The first station to detect an error in a frame which makes the frame invalid (i.e., either an error indicated by the Frame Check Sequence, an invalid frame control field, or an invalid frame length), sets the corresponding E indicator and reports the error to MAC. Once the E indicator has been set, all subsequent stations transmit it as set, but do not report the error. The Address Recognized indicator is set by a station which recognizes the destination address of the frame as its own; all other stations transmit it as received. The Frame Copied indicator is set by a station which recognizes the destination address of the frame as its own and copies the frame into its receive buffer; all other stations transmit the indicator as received.

These status indicators serve a dual purpose. First, they provide a link layer acknowledgement to the sending station. However, since an error may occur in the transfer of the frame from MAC to Logical Link Control, the sending station cannot be assured that its frame was received even if the A and C indicators are set. Also, if the destination address is a group address, the A and C symbols being set only reflects that one of the intended stations received and copied the frame. Second, these indicators help in fault detection and isolation. If a station on the ring recognizes the destination address as its own, but the A flag has already been set, this indicates a duplicate address problem. Comparison of the error counts in the various stations is an aid in detecting and isolating faults, since errors are reported only by the first station to detect the error. Also, a frame with the A indicator set, but not the C indicator, indicates a possible congestion or buffer management problem.

4.1.2. FDDI Station Flags

Each station has three flags, the Address Recognized Flag, the Frame Copied Flag, and the Error Recognized Flag, corresponding to the three mandatory status indicators at the end of each frame. Additional flags in each station record necessary book-keeping information to regulate normal ring operation and information used in ring initialization or ring recovery to indicate which station will issue the new token. All the flags (except one which indicates token class) are cleared when the station receives the first symbol of the starting delimiter of a token or frame and then they are modified according to information in the frame itself.

4.1.3. FDDI Station Counts

The MAC entity in each station maintains some counters "to aid in problem determination and fault location [4]." These counters include Frame_Ct, a count of all frames received, Error_Ct, a count of all frames received with one or more errors that were previously undetected (this count is incremented when the Error Detected status indicator is set by a station), Lost_Ct, a count of all instances when a format error, such as an invalid physical symbol or a symbol that violates the defined format of a frame, is detected when MAC is in the process of receiving a frame or token (MAC strips such frames from the ring, so that they won't be counted by subsequent stations), and Late_Ct, a count of the number of times TRT has expired since the last receipt of a token. Under normal ring operation, Late_Ct should never be greater than 1. The only time a station's Late_Ct can exceed 1 is if its TRT expires while it is transmitting Beacon Frames, because Late_Ct will not get cleared while the transmitter is in this state. The magnitude of Late_Ct is then a measure of the severity of the problem.

Comparison of the counter values within the various stations would be an effective means of fault detection and isolation.

4.2. Performance Analysis

FDDI performance analysis is directed towards detecting and isolating network faults. The Station Management entity in each station is responsible for analyzing status information provided by MAC. MAC reports errors and significant status changes to SMT when they occur. Events to be reported include receipt of a frame that will initiate the ring recovery process, receipt of a frame that signals a disruption of normal ring operation, overflow of station counters, inability to copy a frame addressed to the station, expiration of protocol timers, indication that another station on the ring is responding to this station's address, and any other events as agreed upon between MAC and SMT.

SMT could perform comprehensive data analysis based on these reported statistics. Via peer communication between SMT entities in various stations, the stations could obtain a global picture of network performance. If Station Management detected a potential problem, it could then take whatever action might be appropriate.

4.3. Artificial Traffic Generation

SMT can transmit a frame with a bad frame check sequence, for diagnostic purposes. Other than this, there is no direct provision for artificial traffic generation in FDDI.

5. Configuration Functions

Configuration functions include physical management of stations, address management, station reset, and network reconfiguration. No human intervention should be required for these tasks. The FDDI Station Management, Media Access Control, and Physical Layer Protocol adequately provide for these functions.

5.1. Insertion and Removal

Insertion of a station into the ring is managed automatically by Station Management. This process was discussed in Section 3.5.

Failure of a Class A station which is not equipped with an optical bypass automatically causes reconfiguration of the network. The ring will be wrapped around to physically omit the failed station. Removal of multiple stations from the ring will split the original ring into two or more isolated subrings. A Class B station is either connected to the network or electronically bypassed via a switch in the associated concentrator.

5.2. Address Management

It is SMT's responsibility to assign station addresses and to resolve duplicate address problems. The existence of duplicate addresses is indicated if a station receives a frame with a destination address that it recognizes as its own, but the Address Recognized indicator is already set. This information is reported to SMT, so that SMT can resolve the problem. SMT's actions are not specified at this point, but it might send a frame to the other station to try to resolve the problem. Or, SMT might remove the station from the ring and then reinsert it with a different address.

5.3. Neighbor Notification

It is useful to know the physical topology of the entire network, for both configuration and fault isolation purposes. The global topology of a ring can be easily constructed if each station knows the identity of its upstream neighbor. A neighbor notification procedure is periodically initiated by the active monitor in the IEEE 802.5 token ring. However, the monitor is not essential to the procedure, and neighbor notification could be handled similarly in the FDDI token ring.

5.4. Changing Station Parameters

Each station on the ring contains internal parameters whose setting determines the operational characteristics of that station and of the ring as a whole. Such parameters include timing values used in the negotiation of TTRT, the station's synchronous bandwidth allocation, and parameters identifying events that should be reported to SMT. These parameters can all be assigned by SMT. Change of some parameters might necessitate recovery of the ring, while others might not be so serious. SMT can also clear counters in MAC, so SMT might gather information from MAC and then clear counters on a periodic basis to monitor network performance.

5.5. Reset

Reset of a station's MAC entity can be specifically requested by SMT, or it may be initiated as an automatic by-product of the ring recovery process. MAC reset erases all knowledge of operational parameters within the station. T_Opr is set to the maximum value supported by the station. If this happens to be the value of T_Opr for the rest of the ring, recovery will probably not be initiated. However, if the reset station's T_Opr value is larger than the one used by the rest of the stations, ring recovery is likely to be initiated by another station on the ring. TTRT (and hence, T_Opr) would then be renegotiated, so that all the stations agreed on its value. (See [13] for a more complete discussion of the consequences of station reset.)

5.6. Reconfiguration

SMT manages ring reconfiguration after resource failure or after station insertion as explained in Sections 3.1 and 3.5.

5.7. Synchronous Bandwidth Allocation

This is identified as a function of Station Management in the draft proposed standard, but no details have been provided as yet. Allocation of synchronous bandwidth must be managed so that no more is allocated than the network can support. Also, a means must be provided for dynamic reassignment of these allocations.

6. Comparison to Reliability Mechanisms in Other LAN Architectures

Reliability issues are being addressed in other LAN architectures, as well as FDDI. The IEEE 802 draft proposed LAN standards all refer to network management functions. Draft IEEE standards have been developed for a CSMA/CD bus, 802.3 [8], a token-passing bus, 802.4 [9], and a token ring, 802.5 [10]. A separate document, 802.1, is planned, which is to include a discussion of network management issues as part of its contents.

Considering that the FDDI token ring was designed as a high bandwidth modification of the IEEE 802.5 token ring, it is not surprising that similar reliability mechanisms are included in the two designs. Both specify an interface between the physical layer and network management for insertion of a station into the ring and removal of a station from the ring. Both append status indicators at the end of a frame. Both provide for automatic detection of and recovery from token-related problems, though different timers are required because of the differences in the two access protocols. Both also specify a claim frame process for ring recovery, and a beaconing process for fault isolation. The two interfaces between station (or network) management and medium access control are almost identical. The primary difference in handling ring recovery and fault isolation is the existence of a network monitor in the 802.5 design. Even though every station on the ring has the ability to act as the network monitor, there is only one "active monitor" on the ring at a time and this station has ultimate authority in the ring recovery process. FDDI recovery is completely distributed.

Neither the 802.5 nor the FDDI design for station (or network) management is complete. It is likely that the two will end up having nearly identical functions. The fundamental difference seems to be that 802.5 must include provisions for management of a network monitor, whereas FDDI control is distributed.

Reliability mechanisms provided by FDDI and by the 802.4 token-passing bus are also similar. Ring maintenance functions which are provided for by 802.4 include "ring initialization, lost token recovery, new station addition to logical ring, and general housekeeping of the logical ring [9]." 802.4 specifies an interface between MAC and Station Management, but it is substantially different from the interface provided in FDDI. Most noticeable is the lack of a primitive in 802.4 for MAC to report status to SMT.

Management primitives provided for the 802.3 CSMA/CD bus include a reset command and some access-protocol-related commands. The document also specifies a jabber inhibit capability, i.e., a "self-interrupt capability to inhibit transmit data from reaching the medium [8]," to protect the network from a faulty transmitter that won't stop transmitting. (This is also specified in 802.4.) There is little mention of a network management entity in the document. In fact, it is concluded that for a CSMA/CD access protocol, "no peer management functions are necessary for initiating, terminating, or handling abnormal conditions. Monitoring of on-going activities is done by the carrier sense and collision detect mechanisms [8]."

The goal of the 802 project is to develop a network management entity common to all the architectures specified in the various 802 standards. Conceptually, this is a reasonable goal, since there seems to be no valid reason for the scope of network management capabilities to depend on the particular access protocol, other than to assure the correct functioning of the protocol. In reality, due to the diversity of the various interfaces between MAC and network management as they now exist, development of a common network management may be a difficult task.

Reliability mechanisms are also being included in network implementations that are not being developed according to any of the above IEEE or ANSI standards. Various reliability mechanisms have been included in the design of the Fiber Optic Demonstration System (FODS), a 100 megabit per second fiber-optic star network [20]. This network uses a Carrier Sense Multiple Access with Collision Detection and contention resolution via Time Slot (CSMA/CD/TS) access protocol. FODS was developed by Sperry Space Systems Division in Phoenix, Arizona, under contract to NASA Goddard Space Flight Center, and it is also a candidate for use on the Space Station. FODS includes built-in-test equipment, which tests a backup channel, shuts down a jabbering transmitter, and automatically switches to a backup station upon failure of the primary. The claim is that FODS provides for distributed network control. Network statistics collected at each node "can be used to determine the relative health, efficiency, and loading of the network and individual nodes." An interesting feature of this project is a FODS Exerciser that generates artificial traffic, to determine performance of the network under various conditions, including error situations.

7. Conclusion

The current trend in LAN design, as illustrated by the FDDI effort, is to incorporate network management functions as an extension of OSI layer management. The emphasis is on detecting and correcting errors in the access protocol and on automating physical reconfiguration of the network. In accordance with this philosophy, monitor functions monitor the performance of the access protocol, rather than general performance of the network.

Unfortunately, the standardization of network management functions seems to be a long way away. For even though the importance of including reliability mechanisms in network design is widely accepted, there is no agreement on a specific set of network management tasks that is independent of network topology.

Acknowledgement: The author is indebted to Floyd Ross, Jim Hamstra, and Sunil Joshi, members of the ANSC X3T9.5 committee, for enlightening discussions about FDDI.

References

[1] W. Bux, F. Closs, K. Kuemmerle, H. Keller, and H. Mueller, "Architecture and Design of a Reliable Token-Ring Network," IEEE Journal on Selected Areas in Communications, Vol. SAC-1, No. 5, Nov., 1983, pp.756-765.

[2] "Draft Proposed American National Standard for Local Distributed Data Interfaces," X3T9.5/84-16, Revision 6.0, April 10, 1985.

[3] "FDDI Token Ring Extended Media Access Control Interface," Draft Proposed American National Standard, X3T9.5/84-46, Dec. 13, 1984.

[4] "FDDI Token Ring Media Access Control," Draft Proposed American National Standard, X3T9.5/83-16, Rev 8, Mar. 1, 1985.

[5] "FDDI Token Ring Physical Layer Medium Dependent Standard," Draft Proposed American National Standard, X3T9.5/84-48, Rev 2, April 5, 1985.

[6] "FDDI Token Ring Physical Layer Protocol Standard," Draft Proposed American National Standard, X3T9.5/83-15, Rev 9, April 5, 1985.

[7] "FDDI Token Ring Station Management," Draft Proposed American National Standard, X3T9.5/84-49, Rev 1.2, Feb. 12, 1985.

[8] IEEE Project 802 Local Area Network Standards, "CSMA/CD Access Method and Physical Layer Specifications," Draft IEEE Standard 802.3, Revision D, December, 1982.

[9] IEEE Project 802 Local Area Network Standards, "Token-Passing Bus Access Method and Physical Layer Specifications," Draft IEEE Standard 802.4, Draft D, December, 1982.

[10] IEEE Project 802 Local Area Network Standards, "Token Ring Access Method and Physical Layer Specifications," Draft IEEE Standard 802.5, September 23, 1983.

[11] V. Iyer and S. Joshi, "Streamlining Protocols at High Data Rates," Proceedings of Wescon/84, Session 2, Paper 4.

[12] V. Iyer and S. Joshi, "FDDI's 100 M-bps Protocol Improves on 802.5 Spec's 4 M-bps Limit," EDN, May 2, 1985, pp. 151-160.

[13] M. Johnson, "Proof that Timing Requirements of the FDDI Token Ring Protocol are Satisfied," submitted.

[14] S. Joshi and V. Iyer, "New Standards for Local Networks Push Upper Limits for Lightwave Data," Data Communications, July, 1984, pp. 127-138.

[15] E. Lee and P. Boulton, "The Principles and Performance of Hubnet: A 50 Mbit/s Glass Fiber Local Area Network," IEEE Journal on Selected Areas in Communications, Vol. SAC-1, No. 5, Nov., 1983, pp. 711-720.

[16] T. Phinney, and G. Jelatis, "Error Handling in the IEEE 802 Token-Passing Bus LAN," IEEE Journal on Selected Areas in Communications, Vol. SAC-1, No. 5, Nov. 1983, pp. 784-789.

[17] F. Ross, and R. Moulton, "FDDI Overview - A 100 Megabit per Second Solution," Proceedings of Wescon/84, Session 2, Paper 1.

[18] W. Stallings, *Local Networks: An Introduction*, MacMillan, 1984.

[19] J. Sventek, W. Greiman, M. O'Dell, and A. Jansen, "Token Ring Local Area Networks - A Comparison of Experimental and Theoretical Performance," Proceedings of Computer Networking Symposium, Dec. 13, 1983, IEEE Computer Society Press, 1984, pp. 51-56.

[20] M. Varga, "A Robust 100 MBPS Network for Avionics Application," Proceedings of the National Aerospace and Electronics Conference (NAECON), May, 1985.